arXiv:cond-mat/9908480v3 [cond-mat.stat-mech] 8 Dec 1999

Exact solution of a modified El Farol's bar problem: Efficiency and the role of market impact

Matteo Marsili

Istituto Nazionale per la Fisica della Materia (INFM), Trieste-SISSA Unit, V. Beirut 2-4, Trieste I-34014

Damien Challet

Institut de Physique Théorique, Université de Fribourg, CH-1700

and Riccardo Zecchina

The Abdus Salam International Centre for Theoretical Physics Strada Costiera 11, P.O. Box 586, I-34014 Trieste

Abstract

We discuss a model of heterogeneous, inductive rational agents inspired by the El Farol Bar problem and the Minority Game. As in markets, agents interact through a collective aggregate variable – which plays a role similar to price – whose value is fixed by all of them. Agents follow a simple reinforcement-learning dynamics where the reinforcement, for each of their available strategies, is related to the payoff delivered by that strategy. We derive the exact solution of the model in the "thermodynamic" limit of infinitely many agents using tools of statistical physics of disordered systems. Our results show that the impact of agents on the market price plays a key role: even though price has a weak dependence on the behavior of each individual agent, the collective behavior crucially depends on whether agents account for such dependence or not. Remarkably, if the adaptive behavior of agents accounts even "infinitesimally" for this dependence they can, in a whole range of parameters, reduce global fluctuations by a finite amount. Both global efficiency and individual utility improve with respect to a "price taker" behavior if agents account for their market impact.

⋆ We acknowledge J. Berg, S. Franz and Y.-C. Zhang for discussions and useful suggestions. On the economic side we profited greatly from critical discussions with A. Rustichini and F. Vega-Redondo on learning and dynamic games. This work was partially supported by Swiss National Science Foundation Grant Nr 20-46918.98.

1 Introduction

The El Farol bar problem [1] has become a popular paradigm of complex systems. It describes the situation where N persons have to choose whether to go or not to a bar which is enjoyable only if it is not too crowded. In order to choose, each person forms mental schemes, hypotheses or behavioral rules based on her beliefs and she adopts the most successful one on the basis of past performance. Inductive [1], low [2] or generally bounded rationality based on learning theory [3] is regarded as a more realistic approach to the behavior of real agents in complex strategic situations[4]. This is specially true in contexts involving many heterogeneous agents with limited information, such as the El Farol bar problem. Theoretical advances, beyond numerical simulations, is technically very hard on these problems and it has been regarded as a major step forward in the understanding of complex systems[5].

The minority game [6,7] represents a first step in this direction. It indeed describes a system of interacting agents with inductive rationality which face the problem of finding which of two alternatives shall be chosen by the minority. This problem is quite similar in nature to the El Farol bar problem as the result for each agent depends on what all other agents will do and there is no *a priori* best alternative. These same kind of situations arise generally in systems of many interacting adaptive agents, such as markets[7,8].

Numerical simulations by several authors [6,8–12] have shown that the minority game (MG) displays a remarkably rich emergent collective behavior, which has been qualitatively understood to some extent by approximate schemes [7,13,14]. In this paper, which follows refs. [11,15], we study a generalized minority game and show that a full statistical characterization of its stationary state can be derived analytically in the "thermodynamic" limit of infinitely many agents. Our approach is based on tools and ideas of statistical physics of disordered systems[16].

The minority game, as the El Farol bar problem, allows for a relatively easy definition in words. This may be enough for setting up a computer code to run numerical simulations, but it is clearly insufficient for an analytical approach. Therefore we shall, in the next section, define carefully its mathematical formulation. We shall only discuss briefly its motivation, for which we refer the reader to refs. [6–8]. Even though the behavioral assumptions on which the MG is based may be questionable when applied to financial markets (see sect. 2.4), still we find it convenient to consider and discuss the model as a toy model for a market, in line with refs. [7,8,17]. The relation to markets, at this level, may just be seen as a convenient language to discuss the results in simple terms. This choice reflects our taste and surely more work needs to be done to show the relation of the minority game with real financial markets. We be-

lieve, however, that because of the statistical nature of the collective behavior – which are usually quite robust with respect to microscopic changes – our results may be qualitatively representative of generic systems of agents interacting through a global quantity *via* a minority mechanism, such as markets. The minority game indeed captures the essential interaction between agents beliefs and market fluctuations – how individual beliefs, processing fluctuations, produce fluctuations in their turn. This interaction is usually shortcut in mathematical economy assuming market efficiency, i.e. that prices *instantaneously* react to and incorporate agents beliefs. The motivation underlying the efficient market hypothesis, is, in few words that if there where inefficiencies – or arbitrage opportunities – that would be exploited by speculators in the market and washed out very quickly. Implicitly one is assuming that there is an infinite number of agents in the market who are using very sophisticated strategies which can detect, exploit and eliminate arbitrages very quickly. As "stylized" as it may be, the minority game allows to study how a finite number of heterogeneous agents interact in a complex system such as a market. It allows to ask to what extent this "stylized" market is inefficient and how agents really exploit arbitrage opportunities and to what extent.

After defining the stage game, we shall briefly discuss its Nash equilibria: these are the reference equilibria of deductive rational agents. Finally we shall pass to the repeated game with adaptive agents which follow exponential learning. We show that the key difference between agents playing a Nash equilibrium and agents in the usual minority game is *not* that the first are deductive whereas the latter are inductive. Rather the key issue is whether agents account for their "market impact" or not. By market impact we refer to the fact that the choice of each agent affects aggregate quantities, such as prices. In the minority game [7,8] agents behave as "price takers", i.e. as if their choices did not affect the aggregate. However, due to the minority nature of the interaction, the market impact reduces the "perceived" performance of strategies which agents use in the market with respect to those which which are not used and whose performance is monitored on the basis of a virtual trade (assuming the same price). In order to analyze in detail this issue, we generalize the MG and allow agents to assign an extra reward η to a strategy when it is played. This parameter allows agents to account for their market impact and it plays the same role as the Onsager reaction term, or cavity field, in spin glasses[16].

Our main results are:

(1) We derive a continuum time limit for the dynamics of learning.
(2) We show that this dynamics admits for a Lyapunov function, i.e. a function of all relevant dynamical variables which decreases on all trajectories of the dynamics. This is a very important result since it turns the problem of studying the stationary state of a stochastic dynamical system into that of characterizing the (local) minima of a function. Consider-

ing this function as an Hamiltonian, we can apply the tools of statistical mechanics to solve the problem.

(3) When agents do not consider their impact on the market, as in the minority game ($\eta = 0$), *i)* the stationary state is unique. *ii)* the Lyapunov function is a measure of the asymmetry of the market. In loose words, agents minimize market's predictability. *iii)* When the number of agents exceeds a critical number the market becomes symmetric and unpredictable, with large fluctuations as first observed in [8].

(4) If agents know what is the dependence of the aggregate variable on their behavior they can consider their impact on the market. We refer to this as the *full information* case since agents have full information on how the aggregate would have changed for each of their choices. In this case, *i)* there are exponentially many stationary states. *ii)* these states are Nash equilibria, *iii)* the Lyapunov function measures market's fluctuations, which means that agents cooperate optimally in maximizing global wealth when maximizing their own utility. As a result, fluctuations always decreases as the number of agents increase.

(5) This state is recovered when $\eta = 1$. This means that agents need not have full information in order to reach this optimal state. It is enough that they over-reward the strategy they are currently playing with respect to those they are not playing, by a quantity η.

(6) Any $\eta > 0$ implies an improvement both in individual payoffs – as shown in sect. 8 and in global efficiency with respect to the $\eta = 0$ case.

(7) The most striking result comes when asking how does the collective behavior interpolates between the two quite different limits when changing η from 0 to 1. The result is that when there are few agents the change is mild and continuous – even though there is a phase transition, that is a continuous one (second order). When there are many agents the change happens suddenly and discontinuously as soon as $\eta > 0$. Even an infinitesimal η is enough to reduce market's fluctuations by a finite amount.

These results suggests that the neglect of market impact – which seems an innocent approximation [1] and is usually at the very basis of mathematical economy and finance [2] – plays a very important role in complex systems such as markets.

[1] The impact of each agent on the aggregate is of relative order $1/N$ and it vanishes as $N \to \infty$.

[2] For example, in determining optimal investment strategies or pricing, it is customary to consider prices as just exogenous processes, independent of the trading strategy really adopted.

2 The stage game: strategic structure

2.1 Actions and payoffs

The minority game describes a situation where a large number N of agents have to make one of two opposite actions – such as e.g. "buy" or "sell" – and only those agents who choose the minority action are rewarded. This is similar to the El Farol bar problem, where each one of N agents may either choose to go or not to a bar which is enjoyable only when it is not too crowded. In order to model this situation, let $\mathcal{N} = (1, \ldots, N)$ be the set of agents and let $\mathcal{A} = (-1, +1)$ be the set of the two possible actions. If $a_i \in \mathcal{A}$ is the action of agent $i \in \mathcal{N}$, the payoffs to agent i is given by

$$u_i(a_i, a_{-i}) = -a_i\, A \quad \text{where} \quad A = \sum_{i \in \mathcal{N}} a_i, \tag{1}$$

where $a_{-i} = \{a_j,\, j \neq i\}$ stands for opponents actions. The game rewards the minority group. To see this, note that the total payoff to agents $\sum_i u_i = -A^2$ is always negative. Then the majority of agents, who have $a_i = \operatorname{sign} A$, receives a negative payoff $-|A|$, whereas the minority "wins" a payoff of $|A|$. Eq. (1) can be generalized to $u_i(a_i, a_{-i}) = -a_i\, U(A)$: if the function $U(x)$ is such that $x\, U(x) = -xU(-x) \geq 0$ for all $x \in I\!R$, the game again rewards the minority. The original model[6,7] takes $U(x) = \operatorname{sign} x$, but the collective behavior is qualitatively the same [11] as that of the linear case $U(x) = x$ on which we focus. Note that the "inversion" symmetry $u_i(-a_i, -a_{-i}) = u_i(a_i, a_{-i})$ implies that the two actions are *a priori* equivalent: there cannot be any best actions, because otherwise everybody would do that and loose.

The key issue, clearly, is that of coordination. With respect to coordination games [18, chapt. 6], we remark that agents cannot communicate. If communication were possible, agents would have incentives to stipulate contracts – such as "We toss a coin, if the outcome is head I do $a_{\text{me}} = +1$ and you do $a_{\text{you}} = -1$, and if it is tail we do the other way round". Both players would benefit from this contract because it transforms the negative sum game into a zero sum game for the two players. The contract would then be self-enforcing.

Agents interact only through a *global* or *aggregate* quantity A which is produced by all of them. This type of interaction is typical of market systems [7] ant it is similar to the long-range interaction assumed in mean-field models of statistical physics[16]. Finally note that the El Farol bar problem has a similar structure but with A replaced by $(A - A_0)$ in Eq. (1) where A_0 is related to the bar's comfort level [1,19].

2.2 States and information

Nature can be in one of P *states*, which are labelled by a variable μ which takes integer values $\mu \in \mathcal{P} \equiv \{1, 2, \ldots, P\}$. We assume that P is large and of the same order of N and we define $\alpha \equiv P/N$, which we eventually keep finite in the limit $N \to \infty$. The reason for this particular limit is because, as first shown in ref. [8], the model's behavior only depend on the combination $\alpha = P/N$ in the large N limit. The variable μ encodes all possible information on the state of the environment where agents live, so we shall sometimes call μ "information". μ is drawn from a distribution ϱ^μ on $\mathcal{P}$, independently at any time step. Most results will be presented for the uniform case $\varrho^\mu = 1/P$. In both the El Farol model and in the minority game, μ has a different, more complex definition, on which we shall return later in section 9. In what follows, we shall denote with an over-line $\overline{O} \equiv \sum_{\mu\in\mathcal{P}} \varrho^\mu O^\mu$ the average of a quantity O^μ over μ. For any $\mu \in \mathcal{P}$, payoffs are still given by Eq. (1). Strictly speaking μ is a so called *sun-spot*[20] because the payoffs only depend on the actions of agents. Now, however, the *pure* strategies of each agent may depend on μ. We call $\mathcal{A}^\mathcal{P}$ the set of all such strategies: An element of $\mathcal{A}^\mathcal{P}$ is a function $a : \mu \in \mathcal{P} \to a^\mu \in \mathcal{A}$ or a P dimensional vector with coordinates a^μ, $\forall \mu \in \mathcal{P}$ [3]. There are $|\mathcal{A}^\mathcal{P}| = 2^P$ possible such functions. We call $a_i \in \mathcal{A}^\mathcal{P}$ a possible pure strategy for agent $i \in \mathcal{N}$, with elements $a_i^\mu \in \mathcal{A}$ for all $\mu \in \mathcal{P}$. With this notations, the payoff "matrix" reads $\overline{u_i}(a_i, a_{-i}) = -\overline{a\, A}$ where $A^\mu = \sum_i a_i^\mu$. At this level we have just replicated the game of the previous level P times. Again there cannot be a best strategy $a \in \mathcal{A}^\mathcal{P}$ for the same reasons as before.

2.3 Heterogeneous beliefs and strategies

Now we assume that each agent only restricts her choice on a small subset of S elements of $\mathcal{A}^\mathcal{P}$. We use the vector notation $\vec{a}_i = (a_{1,i}, \ldots, a_{S,i})$ to denote the subset of strategies available to agent i, with elements $a_{s,i} \in \mathcal{A}^\mathcal{P}$. The action of agent i, when the state is μ and she chooses her s^{th} strategy shall then be $a_{s,i}^\mu \in \mathcal{A}$. We shall mainly work on the case where the strategies $a_{s,i}$ are randomly and independently drawn (with replacement) from $\mathcal{A}^\mathcal{P}$. More precisely

$$P(a_{s,i}^\mu = +1) = P(a_{s,i}^\mu = -1) = \frac{1}{2}, \quad \forall i \in \mathcal{N},\ s \in \{1, \ldots, S\},\ \mu \in \mathcal{P}. \tag{2}$$

Note that independence of $\vec{a}_i$ across agents is reasonable because μ is a sunspot and no pre-play communication is possible (agents are assigned their $\vec{a}_i$ before the game starts).

[3] We use the simple letter a without the index $^\mu$ to denote the function.

The utility of agent i, given the value of μ, his choice s_i and the choice of other agents $s_{-i} = \{s_j,\ j \neq i\}$, now becomes

$$u_i^\mu(s_i, s_{-i}) = -a_{s_i,i}^\mu A^\mu \quad \text{with} \quad A^\mu \equiv \sum_{j \in \mathcal{N}} a_{s_j,j}^\mu. \tag{3}$$

The goal of each agent is to maximize his expected payoff over all possible values of μ, which, for agent i, reads

$$\overline{u_i}(s_i, s_{-i}) = -\overline{a_{s_i,i}\, A} \equiv -\sum_{\mu=1}^{P} \varrho^\mu a_{s_i,i}^\mu A^\mu \tag{4}$$

2.4 Motivation

This structure was introduced in refs. [6,7] in order to model *inductive rational* behavior of agents [1]. In few words the argument is the following: If agents have not a completely detailed model of the game they are engaged in, they may think that the value μ has some effect on the game's outcome A, eventually because she believes that other agents will believe the same. This is a self-reinforcing belief because if agents behave differently for different values of μ, the aggregate outcome A will indeed depend on μ, thus confirming agents' beliefs. The game's structure, however, is such that there can be no "commonality of expectations" [1]. This means that, since there is no rational best strategy a_{best} – because otherwise everybody would use that and loose – agents' expectations (or beliefs) are forced to differ.

On the basis of her mental schemes or hypotheses on the situation she faces, an agent may consider a particular strategy more "likely" to predict the "correct" action than another one. More precisely, she may consider that only the S forecasting rules $\vec{a}_i = (a_{1,i} \ldots, a_{S,i})$, out of all the 2^P such rules in $\mathcal{A}^\mathcal{P}$, are "reasonable" or compatible with her beliefs and then restrict her choice to just those. Heterogeneous beliefs are represented by the fact that each agent draws her strategies at random, independently from others using Eq. (2). Alternatively, one may think, following Aumann [21], that the $\vec{a}_i$ are "rules of thumb" that agent i evolved or learned in other contexts and which she applies in this context. We refer the reader to refs. [1,21] for a deeper discussion of this behavior.

It is worth to remark that, in this view, the restriction from 2^P to S strategy is voluntary. It is not difficult to argue that agents have payoff incentives to increase the number of their strategies. In refs. [1,6,7], the state μ is known to agents before they take a decision. So why don't agents take a decision s_i^μ which depends on μ – or even an action a_i^μ which depends on μ? The answer is

that they do not do so because of computational costs one is implicitly building in the model: We are assuming that agents, in complex strategic situations, prefer to simplify decision tasks.

While this motivation may seem reasonable for an issue such as going to a bar or not [1], it may not be appropriate for agents in financial markets as proposed in ref. [7].

On the other hand, this game's structure is justified if we assume that agents do not know the state μ before their choice. Indeed if $\vec{a}_i^{\,\mu}$ where known to agent i before taking her decision, it would be reasonable for her to decide her best action conditional on her *private* information $\vec{a}_i^{\,\mu}$. If μ is not known in advance, one can imagine a situation where agents resort to S "devices" which take actions $a_{s,i}^{\mu}$ for them. The restriction in strategy space in this case is due to some implicit constraints or costs.

Rather then defending the behavioral assumptions of refs. [1,6,7] or pushing further these arguments, we prefer to remain at a generic level. Indeed we believe the model displays collective phenomena which are of a generic relevance, because of their statistical nature. Furthermore, this collective behavior is so rich that, in our opinion, it deserves investigation by its own.

2.5 Notation: mixed strategies and averages

Before coming to the analysis of the game, let us introduce *mixed strategies*. The mixed strategy $\pi_{s,i}$, $s = 1, \ldots, S$ of each agent i is a distribution over her available strategies: $\pi_{s,i}$ is, in other words, the probability with which agent i plays her s^{th} strategy. Again we use a vector notation $\vec{\pi}_i = (\pi_{1,i}, \ldots, \pi_{S,i}) \in \Delta_i$, where Δ_i is the S dimensional simplex of i^{th} agent. We introduce the scalar product $\vec{u} \cdot \vec{v} \equiv \sum_{s=1}^{S} u_s, v_s$ for vectors $\vec{u}, \vec{v} \in I\!R^S$. We also define the norm $|v|^2 \equiv \vec{v} \cdot \vec{v}$ of vectors in $I\!R^S$. Expectations on the mixed strategy of player i reads $E_{\vec{\pi}_i}(\vec{v}) = \vec{\pi}_i \cdot \vec{v}$. We also define the direct product $\Delta^{\mathcal{N}} = \prod_{i \in \mathcal{N}} \Delta_i$ which we shall also call the *phase space*. A point $(\vec{\pi}_1, \ldots, \vec{\pi}_N) \in \Delta^{\mathcal{N}}$ is indeed a possible state of the system. Finally we use the shorthand notation

$$\langle O \rangle = \sum_{s_1=1}^{S} \ldots \sum_{s_N=1}^{S} \pi_{s_1,1} \cdots \pi_{s_N,N} \; O_{s_1,\ldots,s_N} \tag{5}$$

for the expectation on the product measure of mixed strategies over the phase space $\Delta^{\mathcal{N}}$. For example we shall frequently refer to the quantity

$$\langle A^{\mu} \rangle = \sum_{i \in \mathcal{N}} \vec{\pi}_i \cdot \vec{a}_i^{\,\mu}. \tag{6}$$

3 Characterization of collective behavior

As a preliminary to a more detailed discussion, we find it useful to introduce the key quantities which describe the collective behavior. First, as a measure of global efficiency, we take

$$\sigma^2 \equiv \overline{\langle A^2 \rangle} = \sum_{\mu \in \mathcal{P}} \varrho^\mu \left\langle \left(\sum_{i \in \mathcal{N}} a^\mu_{s_i,i} \right)^2 \right\rangle \tag{7}$$

where we remind that $\langle \ldots \rangle$ means expectation over the variables s_i with the corresponding mixed strategy distribution $\pi_{s_i,i}$. Note that $\sigma^2 = -\sum_i \langle \overline{u_i} \rangle$ is just the total loss of agents. A small value of σ^2 implies an efficient coordination among agents.

By construction, the model is symmetric in the sense that, for any μ, no particular sign of A^μ is *a priori* preferred. We shall see, however, that this symmetry can be "broken" resulting in a state where A^μ may take more probably positive than negative values for some μ and *vice-versa* for other values of μ. As a measure of this asymmetry, it is useful to introduce the quantity

$$H \equiv \overline{\langle A \rangle^2} = \sum_{\mu=1}^{P} \varrho^\mu \langle A^\mu \rangle^2 = \sum_{\mu=1}^{P} \varrho^\mu \left(\sum_{i \in \mathcal{N}} \vec{\pi}_i \vec{a}^\mu_i \right)^2 . \tag{8}$$

Note that $H > 0$ implies that there is a *best* strategy $a^\mu_{\text{best}} = -\text{sign}\,\langle A^\mu \rangle$ that could ensures a positive payoff to a new-comer agent. Ideally because if the new-comer really starts playing the game, she will also affect the outcome A^μ. This suggests that H can be regarded as a measure of the exploitable *information content* of the system by an external agent[11].

Both of these quantities are *extensive*, i.e. are proportional to N for $N \to \infty$, and we shall mainly be interested in the finite quantities σ^2/N and H/N. As a statistical characterization of the equilibrium, it is useful to introduce the *self overlap*

$$G = \frac{1}{N} \sum_{i \in \mathcal{N}} |\pi_i|^2 = \frac{1}{N} \sum_{i \in \mathcal{N}} \sum_{s=1}^{S} \pi^2_{s,i} \tag{9}$$

which gives a measure of the average spread of mixed strategies played by agents. If all agents play pure strategies $G = 1$ whereas $G = 1/S$ if $\pi_{s,i} = 1/S$ $\forall s, i$. Therefore $1/G$ is a measure of the "effective" number of strategies that

agents play on average. Note that σ^2 can be written as

$$\frac{\sigma^2}{N} = \frac{H}{N} + 1 - G - \frac{1}{N}\sum_{i\in\mathcal{N}}\sum_{s,s'\neq s}\pi_{s,i}\pi_{s',i}\overline{a_{s,i}a_{s',i}} \cong \frac{H}{N} + 1 - G \qquad (10)$$

where in the last relation we neglected terms which vanish in the limit $N \to \infty$ (because $\overline{a_{s,i}a_{s',i}} \sim P^{-1/2}$ for $s \neq s'$). Eq. (10) means that the loss of agents come either the asymmetry H which they produce or from the stochastic fluctuations of their choices. Indeed if agents play pure strategies, $G = 1$ and the last term vanishes. Put differently, the stochastic fluctuations σ^2 of the market – or volatility – has a systematic contribution H arising from unexploited asymmetries and a stochastic one $1 - G$, which is generated by stochastic choice of agents.

4 Nash Equilibria

Given the payoffs and the choices available to agents we now briefly discuss Nash equilibria of the stage game. The motivation for this section is that, on one hand Nash equilibria shall provide a reference framework for the following discussion. On the other this discussion allows to appreciate the strategic complexity of the problem. For our purposes, Nash equilibria are those states which are stable under payoff incentives, given the choices available to agents. We shall not discuss refinements. We shall remain at a quite simple level without pretending either completeness or rigor.

Given the symmetry of the game, let us first look for symmetric Nash equilibria at the level of actions only (i.e. $P = 1$) with no restriction on strategies ($\vec{a}_i = \mathcal{A}$ for all $i \in \mathcal{N}$). These cannot be in pure actions so let us look for Nash equilibria in mixed actions: Each player either plays $a_i = +1$ with probability π_i or she plays $a_i = -1$ otherwise. It is easy to see that $\pi_i = 1/2$ is the only symmetric Nash equilibrium: No agent has incentives to deviate from the choice $\pi_i = 1/2$ if others stick to it. This state, which we shall call the *random agent state*, is usually taken as a reference state[8,13] and it is characterized by $\sigma^2 = N$ and $H = 0$.

The game has many more Nash equilibria than the symmetric one. For example in pure actions, any state where $|A| = 1$ and N odd (or $|A| = 0$ and N even) is a Nash equilibrium. Indeed, focusing on N odd, agents in the minority (playing $a_i = -A$) would decrease her payoff, switching to the majority side. On the other hand agents in the majority cannot increase their payoff changing from $a_i = A$ to $a_i = -A$ because then, also the majority would change $A \to -A$.

The number $\Omega^{(0)}_{\rm Nash}$ of these Nash equilibria is

$$\Omega^{(0)}_{\rm Nash} = \binom{N}{\frac{N-1}{2}} + \binom{N}{\frac{N+1}{2}}, \qquad \binom{N}{k} \equiv \frac{N!}{k!(n-k)!} \tag{11}$$

which is exponentially large in N. Each of these states is globally optimal, since it has no fluctuation: $\sigma^2 = H = 1$ as compared to $\sigma^2 = N$, $H = 0$ in the symmetric Nash equilibrium.

There are many more Nash equilibria at this simple level [4] . These simple considerations are already enough to appreciate the complexity of the problem. Note in particular that virtually all possible collective behavior, as parametrized by σ^2 and H is possible.

The complexity increases when $P > 1$ and agents have still no restriction on strategies ($\vec{a}_i = \mathcal{A}^{\mathcal{P}}$ for all $i \in \mathcal{N}$) Again we have the symmetric Nash equilibrium – the random agents state – where, for any μ, each agent plays $a = +1$ with probability $1/2$. The number of Nash equilibria in pure actions is huge. Indeed any combination of P pure actions Nash equilibria at the level of actions (one for each value of μ), is a Nash equilibrium. There are therefore $\Omega^{(1)}_{\rm Nash} = \left(\Omega^{(0)}_{\rm Nash}\right)^P$ such equilibria each of them with minimal fluctuations $\sigma^2 = 1$. Again there are many more Nash equilibria.

When the strategies of agents are restricted to the sets $\vec{a}_i$ and S is small (typically finite when $N, P \to \infty$), the problem of identifying Nash equilibria becomes more complex. One way to tackle the problem is to write down the multi-population standard replicator dynamics [22] and then identify Nash equilibria in evolutionarily stable strategies by its stationary and stable points. Again we make no claim of completeness: We just focus on a particular subclass of Nash equilibria [5] – which are evolutionarily stable – which shall play a peculiar role in the following.

The multi-population standard replicator dynamics [22] (RD) reads

$$\frac{d\pi_{s,i}}{dt} = -\pi_{s,i} \sum_{j \in \mathcal{N}, j \neq i} \left[\overline{a_{s,i}(\vec{\pi}_j \cdot \vec{a}_j)} - \overline{(\vec{\pi}_i \cdot \vec{a}_i)(\vec{\pi}_j \cdot \vec{a}_j)} \right]. \tag{12}$$

We observe that $\sigma^2 = \sum_{i,j \neq i} \overline{(\vec{\pi}_i \cdot \vec{a}_i)(\vec{\pi}_j \cdot \vec{a}_j)} + N$ is a Lyapunov function under

[4] Consider e.g. to split the population N into two groups of K agents playing pure actions and $N - K$ playing symmetric mixed actions.

[5] These are *strict* Nash equilibria.

this dynamics. Indeed a little algebra leads to

$$\frac{d\sigma^2}{dt} = -2 \sum_{i \in \mathcal{N}} \sum_{s=1}^{S} \pi_{s,i} \left[\overline{(a_{s,i} - \vec{\pi}_i \cdot \vec{a}_i) \sum_{j \neq i} \vec{\pi}_j \cdot \vec{a}_j} \right]^2 \leq 0. \tag{13}$$

Therefore Nash equilibria are local minima of σ^2 in $\Delta^{\mathcal{N}}$. Furthermore, σ^2 is a linear function of $\pi_{s,i}$ for any i, s, so that $\frac{\partial^2 \sigma^2}{\partial \pi_{s,i}^2} = 0$, $\forall s, i$. Therefore σ^2 is an harmonic function in $\Delta^{\mathcal{N}}$ which implies that the minima are on the boundary of $\Delta^{\mathcal{N}}$. This holds for any subset of variables $\pi_{i,s}$ which therefore implies that minima are located in the corners of the simplex, i.e. Nash equilibria are in pure strategies ($G = 1$). This, in its turn, implies that Nash equilibria have $\sigma^2 \approx H$ by Eq. (10).

A detailed characterization of these Nash equilibria shall be given elsewhere [23]. Here we briefly mention that also in this case Nash equilibria are exponentially many in N. This makes the analytic calculation a step more difficult than the one we shall present later. A simplified approximate calculation (see appendix) gives the following lower bound [6]:

$$\frac{\sigma^2}{N} \geq \begin{cases} \left[1 - \frac{z(S)}{\sqrt{\alpha}}\right]^2 & \text{for } \alpha > z(S)^2 \\ 0 & \text{for } \alpha \leq z(S)^2 \end{cases} \tag{14}$$

where $\alpha = P/N$ and $z(S) = \sqrt{2/\pi} S \int_{-\infty}^{\infty} dz\, e^{-z^2} z [1 - \mathrm{erfc}(z)/2]^{S-1}$ is the expected value of the maximum among S standard random variable (for $S \gg 1$, $z(S) \simeq \sqrt{2 \ln S}$).

Figure 1 shows that the lower bound is already a good approximation to the typical value of σ^2 in the Nash equilibrium, specially for small values of S. Eq. (14) implies that, for fixed S, σ^2 increases with α, which is reasonable because the complexity of information increases and the resources of agents is limited by S. For fixed α, Eq. (14) suggests that σ^2 decreases with S. So if agents are given more resources (larger S), they attain a better equilibrium. Both of these features are confirmed by numerical simulations (see fig. 1).

It is worth to point out that the game specified by the payoffs of Eq. (4), for N and P very large, implies a fantastic computational complexity. Deductive rational agents should be able to master a chain of logical deductions of formidable complexity in order to derive their best response. The efforts required by this strategic situation may well exceed the bounds of memory

[6] Strictly speaking, the meaning of this lower bound is that the probability to observe σ^2 smaller than the lower bound decreases exponentially with N

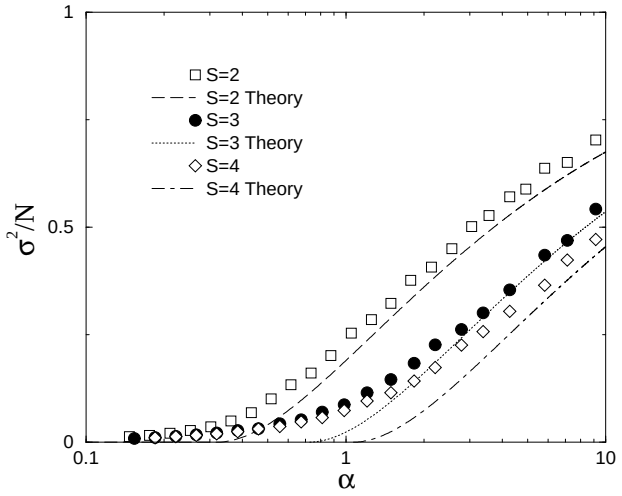

Fig. 1. Global efficiency σ^2/N as a function of α for $S = 2, 3$, and 4 from numerical simulations with $P = 128$ averaged over 100 realizations of $\vec{a}_i$ (symbols) and from the theoretical lower bound (lines).

and computational capabilities of any realistic agent or more simply the resources she is likely to devote to the problem. Furthermore her assumption that everybody else behaves as a rational deductive player becomes more and more unrealistic as N grows large. Finally, even with deductive rational agents there would still be the problem of equilibrium selection which, in this case, involves a huge number of possible equilibria.

5 Repeated game: Learning and inductive rationality

Deductive rationality, as suggested in refs. [1,6,7], is unrealistic in such complex strategic situations [7] and it has to be replaced by inductive rationality. This amounts to assume that agents try to learn what their best choice is from their past performance. We henceforth focus on the repeated game in which agents meet once and again to play the stage game of section 2.3. Different stage games are distinguished by the time label by $t \in I\!N$. For example, $s_i(t)$ denotes the strategy chosen by agent i at time t and $\mu(t)$ the information available at that time.

[7] We do not discuss learning and inductive rationality in simpler cases such as e.g. $P = 1$ and $\vec{a}_i = \mathcal{A}$ for all $i \in \mathcal{N}$. For a discussion of reinforcement learning in El Farol problem at this level see R. Franke (1999).

5.1 Exponential learning:

It is generally accepted that agents follow more likely strategies which have been more successful in the past, which is known[2,24] as the "law of effect". There are several behavioral models implementing the law of effect (see e.g. [2,24]). Here we assume that agents follow an *exponential learning* behavior: Each agent i assigns *scores* $U_{s,i}(t)$ to each of their strategies $s = 1, \ldots, S$ and she plays strategy s with a probability which depends exponentially on its score:

$$\pi_{s,i}(t) = \frac{e^{\Gamma_i U_{s,i}(t)}}{\sum_{s'=1}^{S} e^{\Gamma_i U_{s',i}(t)}} \tag{15}$$

where $\Gamma_i > 0$ is a numerical constant, which may differ for each agent $i \in \mathcal{N}$. This model for discrete choice – called the *Logit model* – has a long tradition in economics [8] [25] and some experimental support (see e.g. [26]). The MG has been originally introduced with $\Gamma_i = \infty$ [6–8] and only recently it has been generalized to $\Gamma_i < \infty$ [27]. Note that $\pi_{s,i}$ here is no more a mixed strategies – which is the object of agents' strategic choice – but rather it encodes a particular behavioral model.

At time $t = 0$ scores are set to some $U_{s,i}(0)$, which encodes *prior beliefs*: e.g. $U_{s,i}(0) > U_{s',i}(0)$ means that agent i considers strategy s *a priori* more likely to be successful than s'.

At later times $t > 0$, agents update the scores of each of their strategies $s = 1, \ldots, S$ in an additive way

$$U_{s,i}(t+1) = U_{s,i}(t) + \Delta U_{s,i}^{\mu(t)}[t, s_i(t), s_{-i}(t)], \tag{16}$$

where the *reinforcement* $\Delta U_{s,i}^{\mu(t)}[t, s_i(t), s_{-i}(t)]$ quantifies the "perceived" success of each of their strategies $s = 1, \ldots, S$ at time t. This generally depends both on the state $\mu(t)$ and on the strategies $s_i(t)$ and $s_{-i}(t)$ played, at that time by agent i and by her opponents.

[8] This model is quite appealing since it satisfies the axiom of *independence from irrelevant alternatives* which states that the relative odds of choices s and s' does not depend on whether another choice s'' is possible or not.

5.2 Naive and sophisticated agents

What is the perceived success $\Delta U_{s,i}$ of a strategy s? The most natural way to quantify the success of a strategy is by the payoff it delivers to the agent if played. This suggests [9] that $\Delta U^{\mu}_{s,i} = u^{\mu}_i[s, s_{-i}]/P$ or

$$U_{s,i}(t+1) = U_{s,i}(t) + u_i^{\mu(t)}[s, s_{-i}(t)]/P. \tag{17}$$

By this equation, however, one assumes that agent i knows what payoff she would have got if she had played any strategy s, including those $s \neq s_i(t)$ which were not used. In other words, agents must have *full information* on the effects (payoffs) of all of their strategies. Furthermore agents take into account the way in which the aggregate quantity $A^{\mu(t)}$ would have changed if they had played strategy $s \neq s_i(t)$. Agents following Eq. (17) are able of sophisticated counter-factual thinking, and are henceforth called *sophisticated* agents [10] . It is worth observing that the score, with the dynamics (17), acquires the meaning of cumulated payoff: $U_{s,i}(t)$ is indeed the payoff agent i would have received (divided by P) if she had always played strategy s against her opponents strategies $s_{-i}(t')$ for all $t' < t$.

In the minority game [11] agents are *naive* [4]: *i)* they eventually have *partial information*, which means that they only know the payoff delivered by the strategy $s_i(t)$ which they actually played. *ii)* they behave as if they were playing against an exogenous signal $A^{\mu(t)}$, rather than $N-1$ other agents. Naive agents neglect their impact on the aggregate A^{μ} and update scores *as if* A^{μ} had not changed if they had used a different strategy. More precisely

$$\Delta U^{\mu(t)}_{s,i}[t, s_i(t), s_{-i}(t)] = -a^{\mu(t)}_{s,i} A^{\mu(t)}(t)/P. \tag{18}$$

Note that Eq. (17) does not depend on the strategy $s_i(t)$ which agent i used, whereas Eq. (18) depends on it because $A^{\mu(t)}(t)$ contains the action $a^{\mu(t)}_{s_i(t),i}$ which agent i actually played. Regarding the MG as a toy market, if the aggregate $A^{\mu(t)}(t)$ plays the role of price, Eq. (18) implies that agents behave as "price takers": They behave as if price did not depend on their actions. By so doing, they simplify considerably the strategic complexity of the context they face. Eq. (18) may be a closer approximation than Eq. (17) to the behavior of real agents in complex strategic situations.

[9] The factor $1/P$ is introduced here and in the following equations for convenience. The reason will become clear in the next section.

[10] This term, as opposed to *naive*, is borrowed from ref. [4].

[11] and probably also in the El Farol problem, ref. [1] is not very clear on this point.

One would naively expect that when $N \to \infty$ the difference between Eqs. (17) and (18) is negligible. Indeed the relative impact of an agent on A^μ is negligible in that limit. Surprisingly we shall see that this is not so and a system of *naive* agents behave quite differently from *sophisticated* agents with full information. In order to study the effect of the impact of agent's choice on the aggregate, we generalize Eq. (18) including a term $+\eta\delta_{s,s_i(t)}/P$. The dynamics of scores then reads

$$U_{s,i}(t+1) = U_{s,i}(t) - a_{s,i}^{\mu(t)} A^{\mu(t)}(t)/P + \eta\delta_{s,s_i(t)}/P. \tag{19}$$

The last term, which is absent ($\eta = 0$) in the original definition of the MG, models the tendency of agents to stick to the strategy they are currently using. Indeed $\eta > 0$ implies that agents reward the strategy they use $s = s_i(t)$ with respect to those they are not currently using $s \neq s_i(t)$. By doing this, agents approximately account for the impact of their actions on the global variable $A^{\mu(t)}(t)$. As we shall see, this term has very deep consequences.

We shall consider these two cases – of *partial information* with *naive* agents and of *full information* – separately below, and see that the collective behavior is remarkably different. Before doing that, we shall first discuss the dynamics of scores in the long run.

6 Continuum time limit and the dynamics in the long run

In this section, we shall first derive a continuum time dynamics for Eq. (16) which captures the long run behavior of the system. Then we shall show that the collective behavior of agents, within this continuum time dynamics, admits a Lyapunov function, i.e. a functions which is minimized along the trajectories of the dynamics of the system. The dynamics therefore converges to the minima of this function. This is a quite important step, since it allows to turn the study of the stationary state of the dynamical model into the study of the local minima of the Lyapunov function. Therefore one can regard the Lyapunov function as the Hamiltonian of a system and resort to the powerful tools of statistical mechanics in order to study the statistical properties of its ground state (global minimum) and eventually of its meta-stable states (local minima). This shall be the subject of the next section.

In order to study the stationary state properties of the system, we need to consider the long time limit of the dynamics of scores, Eq. (16). The key observation, in this respect, is that we expect that $U_{s,i}(t)$ changes significantly and systematically only over time-scales of order $\Delta t \sim P$. Indeed the score of strategies depend on their performance on all the P states μ. In order to

capture the long time dynamics of scores, let us set

$$U_{s,i}(t) = \tilde{U}_{s,i}(\tau), \quad \text{with} \quad \tau = \frac{t}{P}. \tag{20}$$

The dynamics in continuum time of $\tilde{U}_{s,i}$ is obtained iterating Eq. (16) for $\Delta t = P d\tau$ time steps:

$$\frac{\tilde{U}_{s,i}(\tau + d\tau) - \tilde{U}_{s,i}(\tau)}{d\tau} = \frac{1}{P\, d\tau} \sum_{t=P\tau}^{P(\tau+d\tau)-1} \Delta U_{s,i}^{\mu(t)}[t, s_i(t), s_{-i}(t)] \tag{21}$$

We can now take the thermodynamic limit $N, P \to \infty$ keeping $d\tau$ finite. By the law of large numbers, the right hand side converges almost surely to its average value (see later). Here $\Delta U_{s,i}^{\mu(t)}$ is a function of the random variables $\mu(t)$ and $\{s_j(t),\ j \in \mathcal{N}\}$ which are chosen independently at each time step. If the stochastic fluctuations in $\tilde{U}_{s,i}$ are small (see later) also the distribution $\pi_{s,i}(t)$ will be well behaved in the limit $P \to \infty$, specially if Γ_i is small [12]. Indeed defining $\tilde{\pi}_{s,i}(\tau) = \pi_{s,i}(t = P\tau)$, Eq. (15) becomes

$$\tilde{\pi}_{s,i}(\tau) = \frac{e^{\Gamma_i \tilde{U}_{s,i}(\tau)}}{\sum_{s'=1}^{S} e^{\Gamma_i \tilde{U}_{s',i}(\tau)}} \tag{22}$$

which remains meaningful in the limit $P \to \infty$. Equivalently we may say that the distribution $\vec{\pi}_i$ remains approximately constant over $P d\tau$ time steps. Taking the continuum time limit $d\tau \to 0$ in Eq. (21), we find

$$\frac{d\tilde{U}_{s,i}}{d\tau} = \overline{\langle \Delta U_{s,i} \rangle} \tag{23}$$

where averages are taken with respect to the distributions $\vec{\pi}_j$ of strategies and ρ^μ of μ.

It is worth to point out that the order of the two limits – first $P \to \infty$ and then $d\tau \to 0$ – is quite important. Indeed the infinitesimal time interval $d\tau$ corresponds to a very large number $P d\tau$ of time steps, which eventually diverges. This implies that the characteristic time of the system is proportional to P time steps [13].

[12] Interestingly, numerical simulations show that the continuum time approximation works generally also in the limit $\Gamma_i \to \infty$.

[13] In other words, P repetition of the game is the analogous of a "sweep" of the system in a Montecarlo simulation.

The validity of the law of large number can be verified studying the fluctuations of $\Delta U_{s,i}^{\mu(t')}$ around its average $\overline{\langle \Delta U_{s,i} \rangle}$. By routine use of Tchebicev inequality, it is enough to show that

$$\frac{2}{(Pd\tau)^2} \sum_{t,t'=P\tau}^{P(\tau+d\tau)-1} \left\langle \left(\Delta U_{s,i}^{\mu(t)} - \overline{\langle \Delta U_{s,i} \rangle} \right) \left(\Delta U_{s,i}^{\mu(t')} - \overline{\langle \Delta U_{s,i} \rangle} \right) \right\rangle_{\mu,\{s_j\}} \tag{24}$$

vanishes when $N, P \to \infty$. This is indeed the case because the $\Delta U_{s,i}^{\mu(t)}$ depends on the random variables $\mu(t)$ and $\{s_j(t),\ j \in \mathcal{N}\}$, which are drawn independently at each t from their distributions ρ^μ and $\{\vec{\pi}_j,\ j \in \mathcal{N}\}$ respectively. Only the terms with $t = t'$ or $\mu(t) = \mu(t')$ contribute to the average, whereas all other terms vanish because the average factorizes. Only a fraction $\sim 1/P$ of the terms is non-vanishing, which implies that the expression (24) is indeed of order $1/P$ and it vanishes as claimed.

We shall henceforth work with continuum time and drop the tilde over U and π, in order to simplify notations. Combining Eqs. (15) and (23), and with a little algebra, one finds that $\pi_{s,i}$ satisfies the equation

$$\frac{d\pi_{s,i}}{d\tau} = \gamma_i \pi_{s,i} \left[\frac{dU_{s,i}}{d\tau} - \vec{\pi}_i \cdot \frac{d\vec{U}_i}{d\tau} \right] = \Gamma_i \pi_{s,i} \left[\overline{\langle \Delta U_{s,i} \rangle} - \vec{\pi}_i \cdot \overline{\left\langle \Delta \vec{U}_i \right\rangle} \right]. \tag{25}$$

6.1 Lyapunov function for naive agents

The dynamics (19) of naive agents in continuum time is easily derived combining Eqs.(19,25) and a little algebra. It reads

$$\frac{d\pi_{s,i}}{d\tau} = -\Gamma_i \pi_{s,i} \left\{ \sum_{j \in \mathcal{N}} \left[\overline{a_{s,i} \langle a_j \rangle} - \overline{\langle a_i \rangle \langle a_j \rangle} \right] - \eta(\pi_{s,i} - |\pi_i|^2) \right\} \tag{26}$$

with the shorthand $\langle a_j^\mu \rangle = \vec{\pi}_j \cdot \vec{a}_j^\mu$. This is different from RD (Eq. 12). So $\vec{\pi}_i$ does not converge, in general, to a Nash equilibrium. Rather one can show that

$$H_\eta = H - \eta \sum_{i \in \mathcal{N}} |\pi_i|^2 \tag{27}$$

is a Lyapunov function of this dynamics. Indeed, observing that

$$\frac{\partial H_\eta}{\partial \pi_{s,i}} = -2 \frac{dU_{s,i}}{d\tau}$$

and using Eq. (25), one finds that

$$\frac{dH_\eta}{d\tau} = \sum_{i\in\mathcal{N}} \frac{\partial H_\eta}{\partial \vec{\pi}_i} \cdot \frac{d\vec{\pi}_i}{d\tau} = -2 \sum_{i\in\mathcal{N}} \Gamma_i \sum_{s=1}^{S} \pi_{s,i} \left(\frac{dU_{s,i}}{d\tau} - \vec{\pi}_i \cdot \frac{d\vec{U}_i}{d\tau} \right)^2 < 0. \quad (28)$$

The dynamics converges therefore to the minima of H_η.

This equation implies that in the stationary state ($dH_\eta/d\tau = 0$) each of the strategies played by agent i – those with $\pi_{s,i} > 0$ – has the same perceived success $\frac{dU_{s,i}}{d\tau}$ in the long run. Note also that for $\eta = 1$, by Eq. (10), $H_1 \simeq \sigma^2 - N$ which implies that for $\eta = 1$ the stationary state is close to a Nash equilibrium.

We shall come back to the statistical characterization of the stationary state for naive agents in the next section. It is worth to stress, at this point that this result holds for any realization of $a^\mu_{s,i}$. It actually holds for much more general models [28].

6.2 Agents with full information

It is known [24] that exponential learning with full information, for a single agent playing against a stationary stochastic process, converges to rational expectations. If we can regard the opponents of i as a stationary process, this implies that i^{th} strategy converges to the best response. If this happens for all players the system converges to a Nash equilibrium. This is indeed what numerical experiments show (see figure 1).

We can recover this result within the continuum time limit. Indeed, after some algebra one finds that Eq. (25) becomes, in this case

$$\frac{d\pi_{s,i}}{d\tau} = -\Gamma_i \pi_{s,i} \sum_{j\in\mathcal{N},\, j\neq i} \left[\overline{a_{s,i}\langle a_j\rangle} - \overline{\langle a_i\rangle\langle a_j\rangle} \right]. \quad (29)$$

Apart from the factor Γ_i, this coincides with the RD of Eq. (12). Again σ^2 is minimized along the trajectories of Eq. (29): it is easy to check that the time derivative of σ^2 is given by Eq. (13) with an extra factor Γ_i inside the sum on $i \in \mathcal{N}$. We therefore conclude that *with exponential learning and full information agents coordinate on a Nash equilibrium.* Each agent plays, in the long run a pure strategy, i.e. $G = 1$. The Nash equilibrium to which agents converge depends on the initial conditions $U_{s,i}(0)$, i.e. on prior beliefs: Different initial conditions select different Nash equilibria.

7 Statistical mechanics of naive agents

As we have shown, the stationary state of the system is described by the ground state of the Hamiltonian H_η. This can be analyzed using the tools of statistical physics and in particular, the replica method which allows us to deal with quenched disorder (i.e. agents' heterogeneity). The details of the calculation are described in the appendix in some detail. Here we shall just describe the results. We shall consider separately the results for $\eta = 0$, i.e. for the original MG, for which we can derive exact results within a relatively simple calculation. The case $\eta > 0$ requires more complex calculations which shall be the subject of a forthcoming paper [23]. Here, the qualitative understanding for $\eta > 0$ provided by the present approach will be supplemented by numerical simulations.

7.1 $\eta = 0$: The minority game

It is easy to see that for $\eta = 0$, the Hamiltonian $H_0 \equiv H$ is a non-negative quadratic form of the variables $\pi_{s,i}$ and therefore it attains its minimum on a connected subset $\mathcal{M} \in \Delta^{\mathcal{N}}$ [14]. We therefore conclude that *the long run dynamics of this system is described by the minimum of H*. Loosely speaking, in view of our definition of H, we may say that naive agents, in the minority game, minimize the "information content" of the market output $A^{\mu(t)}(t)$.

A complete statistical characterization of the minima of H in the limit $N \to \infty$ with $P/N = \alpha$ finite and $\varrho^\mu = 1/P$, can be obtained from the *replica method* a tool of statistical mechanics devised to deal with disordered systems. An account of this method is given in the appendix together with technical details on the calculation. Here we only discuss the results and their interpretation, We distinguish two regimes separated by a *phase transition* which occurs as $\alpha \to \alpha_c(S) \cong S/2 - 0.6626\ldots$.

7.1.1 Asymmetric phase: $\alpha > \alpha_c$

For $\alpha > \alpha_c$ we find an *asymmetric phase*. Indeed $H > 0$ which means that $\langle A^\mu \rangle \neq 0$ at least for some $\mu \in \mathcal{P}$. The symmetry between the two actions in $\mathcal{A}$ is broken and a *best* strategy $a^\mu_{\rm best} = -\text{sign}\,\langle A^\mu \rangle$ arises in $\mathcal{A}^{\mathcal{P}}$. An $N+1^{\rm st}$ agent who joined the game with this strategy would receive a payoff $\overline{|A|} - \overline{a^2_{\rm best}} =$

[14] Note that, on the contrary, σ^2 is not positive definite and it attains its minima, the Nash equilibria, on a non-connected subset of $\Delta^{\mathcal{N}}$.

$\overline{|A|} - 1$ [15].

The set $\mathcal{M}$ where H attains its minimum consists of a single point, so that, for any initial conditions, the dynamics converges to the same final state in the long run. In other words prior beliefs of agents about their strategies are irrelevant in the long run.

The asymmetry H decreases with decreasing α, which means increasing the number N of agents at fixed P. Naively speaking the asymmetry in $\langle A^\mu \rangle$ is exploited by the adaptive behavior of agents who then reduce it. Indeed agents are more and more selective in their choice, as shown by the fact that G increases as α decreases and the effective number of strategies used $1/G$ decreases. At the same time, as α decreases, the equilibrium becomes more and more "fragile" in the sense that its *susceptibility* to a generic perturbation increases (see the Appendix and ref. [16]).

7.1.2 Phase transition and symmetric phase: $\alpha < \alpha_c$

As $\alpha \to \alpha_c$ the asymmetry vanishes $H \to 0$ and the response of the system to a generic perturbation diverges. This signals a phase transition to the symmetric phase $\alpha < \alpha_c$ where $H = 0$ and any perturbation can change dramatically the equilibrium. The set $\mathcal{M}$ where H attains its minimum is no more a single point, but rather an hyper-plane of a dimension which increases as α decreases. Any point in $\mathcal{M}$ is an equilibrium of Eq. (26) and any displacement along this set can occur freely. In particular, with different initial conditions, the system reaches different points of $\mathcal{M}$. The dynamics (26) indeed converges to the "closest" point on $\mathcal{M}$ which is on its trajectory. In other words, prior beliefs $U_{s,i}(0)$ are relevant for $\alpha < \alpha_c$: With different $U_{s,i}(0)$ the system reaches different equilibria.

7.1.3 Anti-persistence in the symmetric phase

For $\alpha < \alpha_c$ the system is dynamically degenerate: any displacement on the set $\mathcal{M}$ can occur freely. In particular stochastic fluctuations can induce a motion in $\mathcal{M}$. This is what happens for $\Gamma_i \gg 1$ where numerical simulations show the presence of "crowd effects" [8,13,14]. This effect manifests in an increases of σ^2/N as α decreases, which is much faster than what predicted by our theory (see full symbols in fig. 2). This behavior can be traced back to a dynamical anti-persistence [11] resulting from the fact that agents, neglecting their impact on $A(t)$, over-estimate the performance of the strategies they do

[15] The term $-\overline{a^2_{\rm best}}$ is the "market" impact caused by the new agents. It arises because if the strategy $a^\mu_{\rm best}$ where actually played by the $N+1^{\rm th}$ agent, that would also modify $A^\mu \to A^\mu + a^\mu_{\rm best}$. Ref. [17] discusses in greater detail these issues.

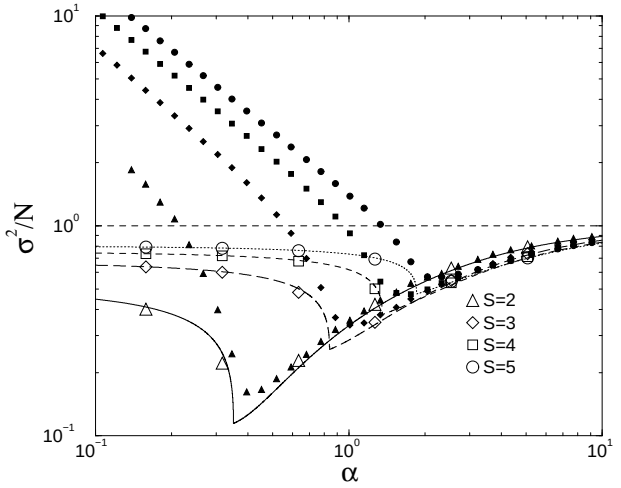

Fig. 2. Global efficiency σ^2/N as a function of α for $S = 2, 3, 4$ and 5 from numerical simulations of the minority game with $N = 101$ agents and $\Gamma_i = \infty$ (averages were taken after $100P$ time steps), averaged over 100 realizations of $\vec{a}_i$ (small full symbols), from numerical minimization of H (large open symbols) and from the theoretical calculation (lines) with $U_{s,i}(0) = 0$.

not play and they keep switching from one to the other. Each time a particular state μ shows up, agents tend to do the opposite action of what they did the last time they saw the same state μ. Therefore the period of this dynamics is of $2P$ time steps [16] [11].

The analytic approach to this effects requires the study of the dynamical solutions of Eq. (19), which go beyond the aims of the present work. We suspect that one should refine our continuum time approach, including eventually the second order time derivative and the effects of fluctuations to some extent. Indeed a periodic motion is usually related to the inertial term ($d^2U/d\tau^2$) of the dynamic equation.

The periodic behavior persists as long as $1/\Gamma_i$ is much smaller than the amplitude of the oscillations of $U_{s,i}(t)$. As Γ_i decreases, Eq. (15) finally smoothes the oscillations in agents choices and the anti-persistent behavior disappears, as indeed observed in [27].

[16] The fact that agents do not realize that they have an impact on the aggregate in this case is probably unrealistic. From their point of view, the same strategy which had a good score of performance when they were not using it, starts performing badly as soon as they use it. This could either be considered a manifestation of *Murphy's law* or the fact that agents rationality is bounded below the level of common sense.

7.1.4 Global efficiency and Tragedy of Commons

As far as global efficiency is concerned, we find that σ^2/N increases with α towards the random agent limit as $1 - \sigma^2/N \sim 1/\alpha$. This is shown in figure 2, which also shows that numerical simulations for finite populations fully confirm our theoretical results. For fixed α, as S increases σ^2/N first decreases moderately as long as $\alpha > \alpha_c(S)$. Then the system enters the symmetric phase ($\alpha < \alpha_c$) because $\alpha_c(S)$ increases, and σ^2/N increases with S towards the random agent limit as $1 - \sigma^2/N \sim 1/S$. We then conclude that allowing agents to have more strategies does not increase global efficiency as in the Nash equilibrium. Rather it pushes the system in the symmetric phase where σ^2 converges to the random agent limit as $S \to \infty$.

As long as the system is in the asymmetric phase, agents have incentives to consider more than S strategies, because that allows them to detect better the asymmetry. However if every agent enlarges the set of strategies she considers, i.e. if S increases, the system enters in the symmetric phase. Then global efficiency starts to decline. This behavior is reminiscent of the *Tragedy of the Commons* [30]: a situation where individual utility maximization by all agents leads to over-exploitation of common resources and poor payoffs.

7.2 Rewarding the played strategy: $\eta > 0$

We expect that with $\eta = 1$ agents behave almost optimally, in the sense that they converge to a stationary state which is close to a Nash equilibrium. Given the difference in the collective behavior of agents in the two cases – which may be appreciated comparing figure 2 for $\eta = 0$ with figure 1 for $\eta = 1$ – it is natural to ask what happens when η changes continuously from 0 to 1.

Figure 3 shows the analytical prediction for the dependence on η of σ^2/N for $S = 2$. These are based on the *replica symmetric ansatz* which is only valid for $\alpha > \alpha_{\rm RSB}(\eta)$, where $\alpha_{\rm RSB}(\eta)$ marks a *replica symmetry breaking* phase transition, which will be discussed elsewhere in detail[23]. Here we just mention that $\alpha_{\rm RSB}(\eta) = 0$ for $\eta < 0$, $\alpha_{\rm RSB}(0) = \alpha_c$ and $\alpha_{\rm RSB}(\eta) \geq 1 - 1/\sqrt{\pi\alpha}$ (for $S = 2$ and) $\eta > 0$. For $\alpha < \alpha_{\rm RSB}(\eta)$ the analytical results derived in the appendix provides an approximate description of the behavior of the system which is however sufficient to appreciate the relevant features.

The most striking consequence of the result in fig. 3 is that the behavior of σ^2/N is quite different for $\alpha > \alpha_c$ and for $\alpha < \alpha_c$. Indeed for large α, σ^2/N changes continuously with η whereas σ^2/N drops discontinuously to zero as $\eta \to 0$ for small α. This feature is reproduced in figure 4 for two characteristic values of α also shown as arrows in fig. 3. We show both the behavior derived from the numerical minimization of H_η and the behavior

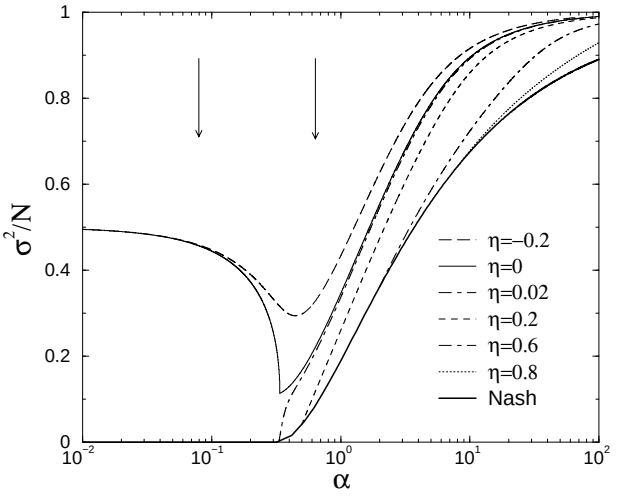

Fig. 3. Theoretical estimate of global efficiency σ^2/N as a function of α for $S = 2$ and several values of η within the *replica symmetric ansatz*.

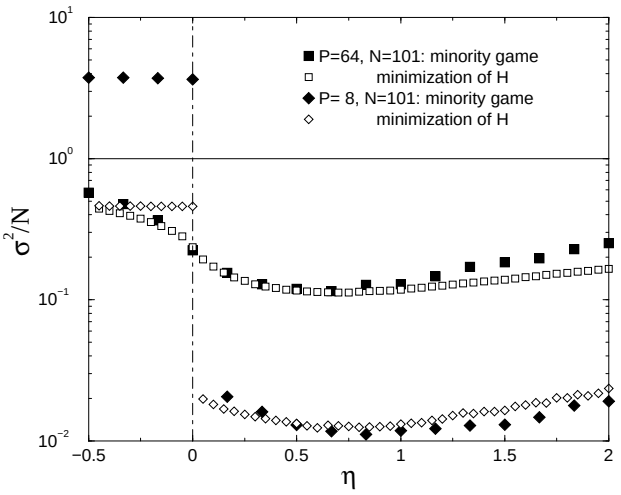

Fig. 4. σ^2/N as a function of η for $S = 2$ and $\alpha \simeq 0.079 < \alpha_c \simeq 0.3374$ and $\alpha \simeq 0.63 > \alpha_c$. Results both of numerical simulations of the minority game and of the numerical minimization of H_η are shown.

of the modified minority game with rewarding. Numerical results agree quite well in an intermediate range of values of η whereas for $\eta > 0.5$ or $\eta < -0.2$ some discrepancy – which we believe is due to finite size effects – is found. This effect can be even more spectacular when anti-persistence effects occur. Indeed the jump of σ^2/N at $\eta = 0$ can be of several orders of magnitude!

The origin of this behavior lies in the dynamic degeneracy of the system for $\alpha < \alpha_c$ and $\eta = 0$. Even an infinitesimal change in η can dramatically alter the nature of the minima of H_η: for negative η there is only one minimum which becomes shallower and shallower as $\eta \to 0^-$. At $\eta = 0$ the minimum is always unique but it is no more point-like. Rather it is a connected set $\mathcal{M}$. An

infinitesimal positive value of η is enough to lift this degeneracy and select only some extreme points of $\mathcal{M}$ as the minima of H_η. The set of minima becomes suddenly disconnected. At fixed $\alpha < \alpha_c$, varying η across the transition H_η changes continuously – with a discontinuity in its first derivative – whereas G and hence σ^2/N change discontinuously with a jump.

The potential implications of this result are quite striking: *rewarding the strategy played more than those which have not been played by a small amount is always advantageous, both individually* (see below) *and globally*. In particular, *an infinitesimal reward is sufficient to avoid crowd effects when α is small and to reduce the fluctuations by a finite amount*.

8 From the agent's viewpoint

Let us consider the behavior of agents in some detail (see also ref. [17] for a related discussion). Our goal is to show that individual payoffs increase when η increases in $[0, 1]$. This means that rewarding the strategy "in vivo", that which has actually been played, is convenient for each agent. We focus on one agent, say i, and assume that others play strategies $s_{-i}(t)$ according to some stationary probability distribution $\vec{\pi}_{-i}$. If $\mu(t)$ is drawn randomly from ρ^μ we can consider $A_{-i}^{\mu(t)}(t) = \sum_{j\neq i} a^\mu_{s_j(t),j}(t)$ as a stationary process. Since we deal with one agent, we shall drop the subscript i in this section. Also we first focus on the case $\eta < 1$ and only after discuss the case $\eta > 1$. In the long run the perceived performance of strategy s is

$$\begin{aligned}\langle \Delta U_s \rangle &= -\overline{a_s \langle A_{-i} \rangle} - \vec{\pi} \cdot \overline{\vec{a} a_s} + \eta \pi_s \\ &\cong -\overline{a_s \langle A_{-i} \rangle} - (1-\eta)\pi_s \end{aligned} \tag{30}$$

where the approximation in Eq. (30) holds for $P \gg 1$ since $\overline{a_{s'} a_s} \sim 1/\sqrt{P}$ for $s' \neq s$. Because of Eq. (25), strategies can either *i)* have $\pi_s > 0$ and $\langle \Delta U_s \rangle = v$ independent of s or *ii)* have $\pi_s = 0$ and $\langle \Delta U_s \rangle < v$. This can be understood by a rather simple argument: Imagine that strategy 1 has $\langle \Delta U_1 \rangle > v$. Then by the very learning dynamics, the agent shall use strategy 1 more frequently than others and hence π_1 shall increase. Because of the last term in Eq. (30) that will decrease the perceived performance $\langle \Delta U_1 \rangle$ of that strategy. On the other hand, if $\langle \Delta U_1 \rangle < v$ the agent shall use it less frequently, hence its $\langle \Delta U_1 \rangle$ shall increase. If $\langle \Delta U_1 \rangle < v$ even when $\pi_1 \to 0$ then the agents will never play that strategy, i.e. $\pi_1 = 0$.

Let $n \leq S$ be the number of strategies with $\pi_s > 0$ and let these be labeled by $s = 1, \ldots, n$, whereas $\pi_k = 0$ for $k > n$. Taking the sum of Eq. (30) on

$s = 1, \ldots, n$ we find

$$v = -\frac{1}{n} \sum_{s=1}^{n} \overline{a_s \langle A_{-i} \rangle} - \frac{1-\eta}{n}$$

where n is fixed by the condition $v > -\overline{a_k \langle A_{-i} \rangle}$ for all $k > n$. Clearly $-\overline{a_s \langle A_{-i} \rangle} > v > -\overline{a_k \langle A_{-i} \rangle}$ for any $s \leq n$ and $k > n$ hence the n strategies which the agent uses are the n more efficient ones. Then Eq. (30) becomes

$$\pi_s = \frac{1}{n} + \frac{1}{1-\eta} \left(\frac{1}{n} \sum_{s'=1}^{n} \overline{a_{s'} \langle A_{-i} \rangle} - \overline{a_s \langle A_{-i} \rangle} \right).$$

Note that strategies with a larger $-\overline{a_s \langle A_{-i} \rangle}$ are played more frequently. The average payoff $\overline{u} = -\vec{\pi} \cdot \overline{\vec{a} \langle A_{-i} \rangle} - 1$ delivered by a learning behavior with parameter η is

$$\overline{u} = -\frac{1}{n} \sum_{s=1}^{n} \overline{a_s \langle A_{-i} \rangle} + \frac{1}{1-\eta} \sum_{s=1}^{n} \left(\overline{a_s \langle A_{-i} \rangle} - \frac{1}{n} \sum_{s'=1}^{n} \overline{a_{s'} \langle A_{-i} \rangle} \right)^2 - 1 \quad (31)$$

which is an increasing function of η for $\eta < 1$. Indeed at fixed n, this is trivially true. With some more algebra, it is easy to check that n is a non-increasing function of η and that $\overline{u}$ increases as n decreases. This means that *for* $\eta < 1$ *average payoffs are non-decreasing functions of* η as claimed.

When $\eta \to 1$ the only possible solution is that with $n = 1$ which means that the agent plays her best response to A_{-i}. For $\eta > 1$ the agent over-weights the performance of her strategies. As a result she sticks to only one of her strategies, i.e. $n = 1$, but that need not be her best one. Without entering in too many details, let us only mention that for $\eta > 1$ the agent plays always one strategy which is dynamically selected by initial conditions and stochastic fluctuations.

9 Exogenous vs endogenous information

In the El Farol problem and in the MG the state $\mu(t)$ is determined by the outcome of past games. In other words $\mu(t)$ is an *endogenous* information which encodes information on the game itself: Agents record which has been the winning action in the last $M = \log_2 P$ games and store this information in the binary representation of the integer μ. This means that μ is updated at

each time as:

$$\mu(t+1) = \mathrm{mod}\left(2\mu(t) + \frac{1+\mathrm{sign}\,A(t)}{2}, P\right), \qquad A(t) = \sum_{i\in\mathcal{N}} a^{\mu(t)}_{s_i(t),i}, \quad (32)$$

Note that $A(t)$ depends on time both through $\mu(t)$ and through the choice $s_i(t)$ of each agent $i \in \mathcal{N}$ at time t. Eq. (32) implies that the dynamics of $\mu(t)$ is defined by the collective behavior of the game itself. Still payoffs do not depend on $\mu(t)$ which is a sun-spot. However agents can coordinate in such a way that some state μ – i.e. some pattern in the time series of $A(t)$ – can occur more frequently than some other and eventually some can never occur. As far as the collective behavior in the stationary state is concerned, the only relevant information of this dynamics is the stationary state distribution ϱ^μ of the process Eq. (32). In the long run, the distribution ϱ^μ is determined by the collective behavior of agents through $A(t)$. Technically, the problem of computing ϱ^μ is related to the diffusion of a particle on the directed graph defined by Eq. (32), where each note $\mu \in \mathcal{P}$ has two outgoing links to nodes $\mathrm{mod}(2\mu, P) + 1$ and $\mathrm{mod}(2\mu + 1, P) + 1$ (and two incoming links). Stochastic processes on this graph, known as De Bruijn graph, are called shift register sequences [31] in computer science.

This differs from the setting we have discussed so far, in which $\mu(t)$ is independently drawn at each time t with $\varrho^\mu = 1/P$. We may call μ *exogenous information* in this case since it can be considered to encode information about an external system, eventually the environment where agents live. This version of the MG has been first introduced by Cavagna [10]. He found that in numerical simulations the collective behavior with exogenous information differs only weakly from that under endogenous information. Having already discussed the results for the exogenous case, let us now consider how these change under endogenous information.

9.1 Naive agents with $\eta \leq 0$ and endogenous information

With endogenous information the system behaves qualitatively in the same way, as first observed in ref. [10] by numerical simulations for $\eta = 0$. This is because the stationary state distribution ϱ^μ of the process $\mu(t)$ – which is induced by the dynamics of agents through Eq. (32) – is almost uniform on $\mathcal{P}$. Actually $\varrho^\mu = 1/P$ for $\alpha \leq \alpha_c$ because of the symmetry of A^μ.

In order to measure the deviation of ϱ^μ from the uniform distribution $\varrho^\mu_{\mathrm{unif}} = 1/P$, we compute the entropy $\Sigma(P) = -\sum_{\mu\in\mathcal{P}} \varrho^\mu \log_P \varrho^\mu$. With the choice of base P for the logarithm $\Sigma(P) = 1$ for $\varrho^\mu = 1/P$ so that $1 - \Sigma(P)$ is a reasonable measure of the deviation of ϱ^μ from a uniform distribution. In

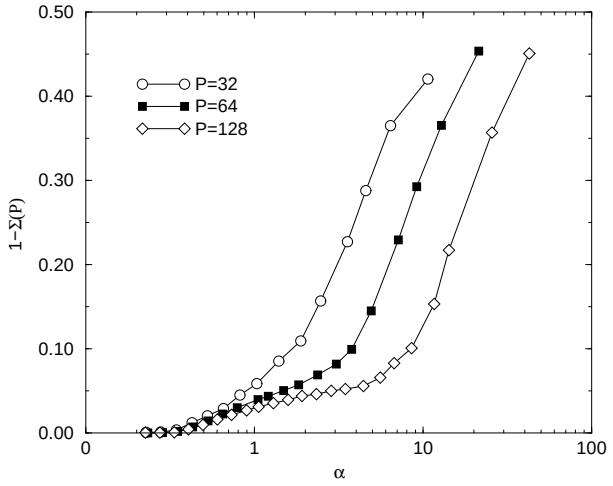

Fig. 5. Deviation of the distribution ϱ^{μ} from the uniform one from numerical simulations with $\eta = 0$.

figure 5, $1 - \Sigma(P)$ is plotted for several values of P as a function of α. While for $\alpha < \alpha_c$ we find $\Sigma(P) = 1$ to a great accuracy, for $\alpha > \alpha_c$ numerical results suggest that $\Sigma(P) \to 1$ as $P = \alpha N \to \infty$. On this basis, we conclude as in ref. [10], that the MG with endogenous information gives the same results as the MG with exogenous information. By a detailed study of the dynamics of the process $\mu(t)$ one can actually give a deeper theoretical foundation to this conclusion and derive analytically this result [32]. We conjecture that, as long as agents choice remains stochastic ($G < 1$), as for $\eta \leq 0$, the dynamics in $\mu(t)$ is ergodic in $\mathcal{P}$. Note indeed that, for any μ, $A(t)$ has stochastic fluctuations around its average value $\langle A^{\mu} \rangle$ which are of the same order of magnitude of the average itself.

9.2 $\eta > 0$ and agents with full endogenous information

A qualitatively different situation arises when $\eta > 0$ and in particular when agents have full information. We shall mainly discuss the latter case which corresponds to $\eta = 1$ and then discuss briefly the generic $\eta > 0$ case. The key observation is that agents for $\eta = 1$ play pure strategies ($G = 1$). This means that agents behave the same whenever $\mu(t) = \mu$ and accordingly $A(t)$ shall always take the same value A^{μ} each time $\mu(t) = \mu$. This implies that the dynamics of Eq. (32) for $\mu(t)$ becomes deterministic. More precisely it locks into a periodic orbit $\mu(t + T) = \mu(t)$ with some period T. Only the values of μ into this orbit shall occur in the long run, whereas all other values of μ shall never occur. This means that agents strongly influence the time series of $\mu(t)$ and hence of the aggregate $A(t)$. Most remarkably, in doing so, they achieve a much better coordination with respect to the exogenous information

case because they reduce the parameter $\alpha = P/N$ to $\tilde{\alpha} = T/N$ [17]. Numerical studies show that $T \propto \sqrt{P}$ so that, in the limit $N \to \infty$ with $P/N = \alpha$ finite, $\tilde{\alpha} \propto 1/\sqrt{N}$ and also σ^2/N takes a very small value [18].

The same behavior shall be expected for all $\eta > 0$ such that $G = 1$. Therefore we expect that for each α there shall be a particular value $\eta_{\rm EB}(\alpha)$ beyond which ergodicity of the dynamics of $\mu(t)$ in $\mathcal{P}$ breaks down. For values of $\eta > \eta_{\rm EB}(\alpha)$ we expect the dynamics of μ to lock into periodic orbits causing a reduction of σ^2 similar to that discussed above.

10 Discussion

There is a growing literature on learning which addresses the issue of which learning procedure (modeling inductive rationality) may eventually lead to deductive rational outcomes[3,24]. The choice made in the El Farol bar problem and in the MG – which is exponential learning (Eq. 15) eventually with $\Gamma_i \to \infty$ – is one of these as we have shown. What leads to equilibria different from Nash equilibria is the fact that agents *i)* have not full information on the effects of their strategies and *ii)* that they neglect or do not account properly for their impact on the aggregate in the evaluation of their strategies.

In particular, for $\eta = 0$, this new equilibrium, which we call *naive agents equilibrium* (NAE) differs substantially from a Nash equilibrium (NE), because:

(1) In a NE global efficiency always increases as the number N of agents increases (with P fixed), whereas in the NAE it only decreases as far as $N < P/\alpha_c$ and then it increases in a way which depends on initial conditions (prior beliefs) and on the parameters Γ_i.
(2) There are (exponentially) many *disconnected* Nash equilibria which are selected by initial conditions, i.e. by prior beliefs. For $\alpha > \alpha_c$ there is a unique NAE and, for all initial conditions, the system converges to it. For $\alpha < \alpha_c$ there is a continuum of NAE, but they are all connected.
(3) Global efficiency (σ^2), for fixed α, always decreases as agents resources (S) increase, and it eventually converges to perfect optimization for $S \to \infty$. In the NAE efficiency only mildly improves with S in the asymmetric phase. But increasing S also increases $\alpha_c(S)$ and when $\alpha_c(S) > \alpha$ the system enters into the symmetric phase where σ^2 increases with S towards

[17] Note that the values of μ which occur in the long run are sampled uniformly.
[18] The result $T \propto \sqrt{P}$ is what one would obtain on a random directed graph with P vertex each with two outgoing links. This is a reasonable approximation because the dynamics locks into periodic loops of $T \propto \sqrt{P}$ vertices where the peculiar structure of De Bruijn graphs (see Eq. (32)) does not play a significant role.

the random agent limit (this occurs for Γ_i small [27]).

(4) In the NE agents play *pure* strategies, i.e. $\vec{\pi}_i$ is a singleton $\forall i \in \mathcal{N}$. Indeed, for fixed opponent strategies s_{-i}, each agent typically [19] has a pure strategy s_i – the best response – which is superior to others. In the NAE, agents mix this strategy with others because, neglecting their impact on $A(t)$, they over-estimate the performance of the strategies they do not play. Playing a pure strategy reduces its perceived performance and this is why agents mix strategies in the NAE (see also ref. [17]). The probability $\pi_{s,i}$ in the NAE is such that the perceived performance of all strategies which are played ($\pi_{s,i} > 0$) is the same.

(5) A consequence of the previous point is that, the origin of information is quite important for inductive agents with full information, while it is irrelevant in the NAE [10]. Inductive agents with full information lead, under *endogenous* information, to a deterministic dynamics of $\mu(t)$ and only a small subset of informations μ is ever visited. This in turn leads to a much more efficient coordination. Naive agents with endogenous information, on the other hand, induce a dynamics on $\mu(t)$ which is "ergodic", i.e. which visits each information μ with nearly the same frequency ϱ^{μ}. Therefore the collective behavior is the same as that of the NAE with exogenous information.

For intermediate values of η the collective behavior of agents interpolates between these two situations in a continuous way for $\alpha > \alpha_c$ or in a discontinuous way for $\alpha < \alpha_c$.

We expect that rewarding the strategy which is played with respect to those which are not played should be advantageous both individually and globally in more general situations where agents interact through a global variable *via* a minority-like mechanism. Indeed one can show that the qualitative picture we have described remains the same if we allow for heterogeneity of various sorts such as allowing for a dependence on i of S and η, or changing the distribution in Eq. (2) to a generic $P_i(a)$ for $a \in I\!R$ [28] (see the appendix).

Our results clearly allow for several extensions, as those of ref. [17]. It also suggests a theoretical approach to the El Farol problem [1] [20]. We expect

[19] Typically here means that we disregard unlikely realizations where two pure strategies happen to yield the same payoffs.

[20] The key problem lies in the parametrization of forecasting rules: agents in the El Farol problem consider the record of past attendance to the bar – i.e. $A(t')$ for $t' < t$ – whereas in the minority game agents only consider the record of the sign of $A(t')$. Focusing on the last M games, there can be N^M possible records in the El Farol problem, instead of 2^M. This causes no problem in principle since one can take $P = N^M$. In practice however forecasting rules have to be "reasonable": For example a randomly drawn rule can easily predict a different outcome if the attendance of one of the past weeks just changes by one unit. Some sort of continuity

that the distinction between inductive agents with full information and naive agents to be of key importance also in the El Farol problem and we believe that, eventually, an analytical solution in the limit $N \to \infty$ is possible along the same lines followed here.

A Appendix: Replica calculation for the MG

Our goal is to compute and characterize the minimum of $H_\eta = \overline{\langle A \rangle^2} - \eta N G$, with G given by Eq. (9), in $\Delta^{\mathcal{N}} = \{\vec{\pi}_i,\ i \in \mathcal{N}\}$. Considering H_η as an Hamiltonian of a statistical mechanic's system, this can be done analyzing the zero temperature limit. First we build the partition function

$$Z(\beta) = \mathrm{Tr}_\pi e^{-\beta H_\eta\{\pi\}}, \tag{A.1}$$

where β is the inverse temperature and Tr_π stands for an integral on $\Delta^{\mathcal{N}}$ (we call simply π an element of $\Delta^{\mathcal{N}}$). The quantity of interest is then

$$\min_{\pi \in \Delta^{\mathcal{N}}} H_\eta\{\pi\} = -\lim_{\beta \to \infty} \beta^{-1} \ln Z(\beta). \tag{A.2}$$

This in principle depends on the specific realization $a^\mu_{s,i}$ of rules chosen by agents. In practice however, to leading order in N, all realizations of $a^\mu_{s,i}$ yield the same limit, which then coincides with the average of $\min_{\pi \in \Delta^{\mathcal{N}}} H_\eta\{\pi\}$ over $a^\mu_{s,i}$. The average of $\ln Z$ over the a's, which we denote by $\langle \ldots \rangle_a$, is reduced to that of moments of Z using the replica trick[16]:

$$\langle \ln Z \rangle_a = \lim_{n \to 0} \frac{1}{n} \ln \langle Z^n \rangle_a \tag{A.3}$$

With integer n the calculation of $\langle Z^n \rangle_a$ amounts to study n *replicas* of the the same system with the same realization of $a^\mu_{s,i}$. To do this we introduce a set of dynamical variables $\pi_a \equiv \{\pi_{s,i,a}\}$ for each replica, which are labeled by the additional index $a = 1, \ldots, n$. Each replica has its corresponding Hamiltonian, which we write as $H^a_\eta\{\pi_a\} = \overline{A^2_a} - \eta N G_{a,a}$ where $A^\mu_a = \sum_{i \in \mathcal{N}} \vec{\pi}_{i,a} \cdot \vec{a}_i$ and $N G_{a,a} = \sum_i |\pi_{i,a}|^2$ (the reason for this notation shall become clear later). The set of all dynamical variables for all replicas is the direct product $\Delta^{\mathcal{N}n}$ of n phase spaces $\Delta^{\mathcal{N}}$. In order to compute the limit $n \to 0$ in Eq. (A.3) one appeals to analytic continuation of $\langle Z^n \rangle_a$ for real n. We give here the details of

in strategies should be introduced so that "similar" histories (=information μ) lead to similar forecasts. Even though it is not clear how to translate this requirement mathematically (one way could be that followed in ref. [33]) it is clear that it suggests that the number of relevant informations P is much less than N^M.

the calculation in our specific case. More details on the nature of the method can be found in ref. [16]. We write

$$\begin{aligned}\langle Z^n\rangle &= \mathrm{Tr}_\pi \prod_{a=1}^{n}\prod_{\mu\in\mathcal{P}}\left\langle e^{-\beta\varrho^\mu[(A_a^\mu)^2-\eta N G_{a,a}]}\right\rangle_a \\ &= \mathrm{Tr}_\pi \prod_{\mu\in\mathcal{P}}\prod_{a=1}^{n} E_{z_a^\mu}\left\langle e^{-i\sqrt{2\beta\varrho^\mu}z_a^\mu A_a^\mu}\right\rangle_a e^{\beta\eta N G_{a,a}} \end{aligned} \tag{A.4}$$

where $E_z[\ldots]$ stands for the expectation over the Gaussian variable (unit variance an zero mean) z and we have introduced one such variable z_a^μ for each a and μ, using the identity $E_z[e^{-ixz}] = e^{-x^2/2}$. In addition we used the shorthand Tr_π for the integral over $\Delta^{\mathcal{N}n}$. The average over $a_{s,i}^\mu$ now factorizes

$$\prod_{a=1}^{n}\left\langle e^{-i\sqrt{2\beta\varrho^\mu}z_a^\mu A_a^\mu}\right\rangle_a = \prod_{i\in\mathcal{N}}\prod_{s=1}^{S}\left\langle e^{-i\sqrt{2\beta\varrho^\mu}(\sum_a z_a^\mu \pi_{s,i}^a)a_{s,i}^\mu}\right\rangle_a =$$

and we can explicitly compute it using the distribution (2). This gives

$$= \prod_{i\in\mathcal{N}}\prod_{s=1}^{S}\cos\left[\sqrt{2\beta\varrho^\mu}\sum_{a=1}^{n} z_a^\mu \pi_{s,i}^a\right] \simeq \prod_{i\in\mathcal{N}}\exp\left[-\beta\varrho^\mu\sum_{a,b=1}^{n} z_a^\mu z_b^\mu \sum_{i\in\mathcal{N}}\vec{\pi}_i^a\cdot\vec{\pi}_i^b\right].$$

In the last passage we used the relation $\cos x \simeq e^{-x^2/2}$ which is correct to order x^2 in a power expansion. This is justified as long as $\varrho^\mu \to 0$ as $P = \alpha N \to \infty$ for each $\mu \in \mathcal{P}$. Note that, because of this reason, we would have got the same result for any generic distribution $P_i(a)$ of $a_{s,i}^\mu$ such that $\langle a\rangle = 0$ and $\langle a^2\rangle = 1$. This allows us to understand why models with continuum strategies $a_{s,i}^\mu \in I\!R$, such as the one proposed in ref. [27], yield the same results as the one with binary strategies, which we are discussing here. Before going back to Eq. (A.4), we introduce the matrices $\hat{G} \equiv \{G_{a,b},\ a,b = 1,\ldots,n\}$ and $\hat{r} \equiv \{r_{a,b},\ a,b = 1,\ldots,n\}$ through the identities

$$1 = \int dG_{a,b}\,\delta\left(G_{a,b} - \frac{1}{N}\sum_{i\in\mathcal{N}}\vec{\pi}_i^a\cdot\vec{\pi}_i^b\right) \propto \int dr_{a,b}dG_{a,b}\,e^{\frac{\alpha\beta^2 r_{a,b}}{2}(\sum_i \vec{\pi}_i^a\cdot\vec{\pi}_i^b - NG_{a,b})}$$

for all $a \geq b$, where $\delta(x)$ is Dirac's delta function and we used its integral representation. The only part depending on the $\pi_{s,i}^a$ in $\langle Z^n\rangle$ is $e^{\alpha\beta^2\sum_{a\geq b} r_{a,b}\sum_i \vec{\pi}_i^a\cdot\vec{\pi}_i^b/2}$. This can be factorized in the agent's index i and so the integral Tr_π on $\Delta^{\mathcal{N}n}$ can be factorized into N integrals over Δ^n (=the direct product of the simplexes of the n replicas of the same agent's mixed strategies). With this we can write

$$\langle Z^n\rangle = \int dr_{a,b}dG_{a,b}\,e^{-\beta n N F_\beta(\hat{G},\hat{r})} \tag{A.5}$$

where, specializing to the case $\varrho^\mu = 1/P$ [21] ,

$$F_\beta(\hat{G},\hat{r}) = \frac{\alpha}{2n\beta}\ln\det\left[\hat{I}+\frac{2\beta}{\alpha}\hat{G}\right] + \frac{\alpha\beta}{2n}\sum_{a,b} r_{a,b}G_{a,b} \\ -\frac{1}{n\beta}\ln \mathrm{Tr}_{\pi\in\Delta^n}\exp\left[\frac{\alpha\beta^2}{2}\sum_{a,b} r_{a,b}\vec{\pi}^a\vec{\pi}^b\right] - \eta\sum_a G_{a,a}, \qquad \text{(A.6)}$$

where $\hat{I}$ is the identity matrix. The first term arises from the expectation over z_a^μ. This factorizes for each μ and one is left with a Gaussian integral over $\vec{z}\in I\!R^n$. The second and the third terms arise from the integral representation of the delta functions [22] .

The key point is that, in the limit $N\to\infty$ the integral over the matrices $\hat{r}$ and $\hat{G}$ in Eq. (A.5) are dominated by their saddle point value, i.e. by the values of $r_{a,b}$ and $G_{a,b}$ for which F attains its minimum value [23] . One should then study the first order conditions $\partial F/\partial r_{a,b}=0$ and $\partial F/\partial G_{a,b}=0$ for all a,b. Here we focus on the *replica symmetric* approximation where we assume that the matrices for which F attains its extreme have the form

$$G_{a,b} = g + (G-g)\delta_{a,b}, \qquad r_{a,b} = r + (R-r)\delta_{a,b}. \qquad \text{(A.7)}$$

This ansatz is correct for $\eta\le 0$ and for $\eta>0$ and α large enough[23]. The reason for this is that H_η is a non-negative definite quadratic form in $\Delta^{\mathcal{N}}$. Hence it has a very simple *energy landscape*, characterized by a single *valley*. Taking the limit $n\to 0$, Eq. (A.3) then gives

$$F_\beta^{(RS)}(Q,q,R,r) = \frac{\alpha g}{\alpha+2\beta(G-g)} + \frac{\alpha}{2\beta}\ln\left[1+\frac{2\beta(G-g)}{\alpha}\right] - \eta G \\ +\frac{\alpha\beta}{2}(RG-rg) - \frac{1}{\beta}E_{\vec{z}}\{\ln \mathrm{Tr}_\pi \exp\left[-\beta V_{\vec{z}}(\vec{\pi})\right]\} \qquad \text{(A.8)}$$

where Tr_π is now the integral over the simplex Δ of a single agent's mixed strategies and we defined, for convenience, the potential $V_{\vec{z}}(\vec{\pi}) = \sqrt{\alpha r}\,\vec{z}\cdot\vec{\pi} - \frac{\alpha}{2}\beta(R-r)|\pi|^2$. The parameters g, G, r and R are fixed by the first order conditions $\partial F_\beta^{(RS)}/\partial g = 0$, $\partial F_\beta^{(RS)}/\partial G = 0$, $\partial F_\beta^{(RS)}/\partial r = 0$ and $\partial F_\beta^{(RS)}/\partial R =$

[21] A generic distribution ϱ^μ can also be handled, though with heavier notations.
[22] For simplicity we have also done the transformation $r_{a,b}\to r_{a,b}/2$ for $a\neq b$ so that $\sum_{a\ge b}\to\sum_{a,b}$.
[23] Note that, by Eq. (A.2), we shall also be interested in the limit of $\beta\to\infty$ in the end!

0. These equations, finally, have to be studied in the limit $\beta \to \infty$, where one recovers the minimum of H_η by Eq. (A.2), i.e.

$$\lim_{N\to\infty} \min_{\pi\in\Delta^{\mathcal{N}}} \frac{H_\eta\{\pi\}}{N} = -\lim_{\beta\to\infty} \frac{1}{\beta N} \langle \ln Z(\beta) \rangle = \lim_{\beta\to\infty} F_\beta^{(RS)}(Q,q,R,r)\Big|_{\rm sp}$$

where the subscript sp means that we compute the function $F_\beta^{(RS)}$ at the saddle point values of Q, q, R and r.

It is convenient to define the parameters

$$\chi = \frac{2\beta(G-g)}{\alpha}, \qquad y = \frac{\sqrt{g/\alpha}}{1+\eta(1+\chi)} \tag{A.9}$$

In the limit $\beta \to \infty$, we first look for solutions where $g \to G$ and χ, which we call *susceptibility*, remains finite. This implies that two replicas of the same system converge in the long run to the same stationary state. Using the saddle point equations, and $g = G$, we can rewrite

$$V_{\vec{z}}(\vec{\pi}) = \frac{2y\,\vec{z}\cdot\vec{\pi} + \pi^2}{1+\chi}, \qquad \beta \to \infty \tag{A.10}$$

The last term in Eq. (A.8) is dominated by the mixed strategy $\vec{\pi}^*(\vec{z})$ which is the solution of

$$\vec{\pi}^*(\vec{z}) = \arg\min_{\pi\in\Delta} V_{\vec{z}}(\vec{\pi}). \tag{A.11}$$

We find that $G = g = E_{\vec{z}}[\vec{\pi}^*(\vec{z})]$, which is then a function of y only $G \equiv G(y)$. Upon defining $\zeta(y) = E_{\vec{z}}[\vec{z}\cdot\vec{\pi}^*(\vec{z})]$, we find

$$\chi(y) = -\frac{\zeta(y)}{\sqrt{\alpha G(y)} + \zeta(y)}$$

The second of Eqs. (A.9) becomes an equation for y as a function of α which has two implicit solutions. These can be expressed as explicit solutions for α as a function of y and η:

$$\alpha = \frac{1}{G}\left[\frac{G - y\zeta \pm \sqrt{(G+y\zeta)^2 - 4\eta y\zeta G}}{2(1-\eta)y}\right]^2 \tag{A.12}$$

A.1 $\eta \leq 0$

The solutions of Eq. (A.12), for $\eta < 0$ describe the two branches $\alpha < \alpha_c$ and $\alpha > \alpha_c$. In particular for $\eta \to 0^-$ these solutions become

$$\alpha y^2 = G(y), \qquad \alpha G(y) = \zeta^2(y). \tag{A.13}$$

Let us discuss first the case $\eta \to 0^-$: The *free energy* per agent is

$$\lim_{N\to\infty} \frac{H}{N} = -\lim_{N\to\infty} \frac{\langle \ln Z(\beta) \rangle_a}{N} = \frac{G}{(1+\chi)^2} \tag{A.14}$$

These equations are transcendental and we could not find an explicit solution for generic S. Nevertheless, they represent a great simplification with respect to the original problem. The main technical difficulty lies in the evaluation of the functions $G(y) = E_{\vec{z}}[|\pi^*(\vec{z})|^2]$ and $\zeta(y) = E_{\vec{z}}[\vec{z} \cdot \vec{\pi}^*(\vec{z})]$, which can be computed numerically to any desired accuracy $\forall S$.

The first of Eqs. (A.13) gives the $\alpha > \alpha_c$ phase. This solution has $\chi > 0$ finite and $H > 0$ non-zero. As α decreases χ increases and it diverges as $|\alpha - \alpha_c|^{-1}$ when $\alpha \to \alpha_c^+$. In this limit Eq. (A.14) implies that $H \sim |\alpha - \alpha_c|^2$ vanishes. The critical point $\alpha_c = \alpha(y_c)$ is obtained imposing $\chi = \infty$, which gives $G(y_c) = -y_c \zeta(y_c)$. By the numerical evaluation of the functions $G(y)$ and $\zeta(y)$, we find

$$\alpha_c(S) \cong \alpha_c(2) + \frac{S-2}{2} \tag{A.15}$$

to a high degree of accuracy. It might be that this equation is exact but we could not prove it. An interesting relation for $\alpha_c(S)$ can be derived by algebraic considerations: Note that for each $\pi_{s,i} > 0$ the equation

$$\frac{\partial H}{\partial \pi_{s,i}} = 2 \sum_{j,s'} \overline{a_{s,i} a_{s',j}} \pi_{s',j} = 0 \tag{A.16}$$

must hold. This is a set of linear equations in the variables $\pi_{s,i} > 0$. The $NS \times NS$ matrix $\overline{a_{s,i} a_{s',j}}$ is built with P dimensional vectors $a^\mu_{s,i}$ and therefore has at most rank P. In other words there are only P independent equations (A.16). In addition there are N normalization conditions on $\pi_{s,i}$. The system becomes dynamically degenerate when the number of free variables $\pi_{s,i}$ becomes bigger than the number $P + N$ of independent equations and, exactly at α_c the two

are equal. Dividing this condition by N gives the desired equation

$$\sum_{s=1}^{S} E_{\vec{z}}\{\theta[\pi_s^*(\vec{z})]\} = \alpha_c(S) + 1. \tag{A.17}$$

The left hand side is the average number of strategies used by agents (called n in section 8). Note that this equation implies that $\alpha_c(S)$ cannot grow faster than linear in S. Also $\alpha_c(S) \propto S/2$ imply that agents use on average $1/2$ of their strategies at α_c.

The second of Eqs. (A.13) gives the $\alpha < \alpha_c$ phase. Note indeed that with this choice $\chi \simeq -1/\eta \to \infty$ and $H \sim \eta^2 \to 0$ as $\eta \to 0^-$. At odds with the solution for $\alpha > \alpha_c$, this equation only arises if $\eta < 0$ and in the limit $\eta \to 0^-$. With $\eta = 0$ the saddle point equations have only a solution with $G > g$ in the limit $\beta \to \infty$. This is because for $\alpha < \alpha_c$ the set $\mathcal{M}$ where $H = 0$ is not a single point, but rather a connected set. The replica method with $\eta = 0$ takes an average on all the set $\mathcal{M}$ and so it gives results which are not representative of a particular system [24]. In order to select a single point in $\mathcal{M}$ one may consider the limit $\eta \to 0^-$. Since the term $-\eta\, N\, G$ in H_η breaks the degeneracy of equilibria for $\eta = 0$, the limit $\eta \to 0^-$ selects the equilibrium which is closest to the random initial condition $\pi_{s,i}(0) = 1/S$ for all $i \in \mathcal{N}$ and $s = 1, \ldots, S$. This describes the stationary state of a system of agents with no prior beliefs ($U_{s,i}(t = 0) = 0$, $\forall s, i$).

In both phases, once the saddle point equations are solved, one can derive the full statistical characterization of the system. For instance the fraction of agents playing a strategies in a neighborhood $d\vec{\pi}$ of $\vec{\pi}$ is given by $p(\vec{\pi})d\vec{\pi} = E_{\vec{z}}[\delta(\vec{\pi}^*(\vec{z}) - \vec{\pi})]d\vec{\pi}$.

A.2 $\eta > 0$

Let us for simplicity consider the simpler case $S = 2$. The solution with $G = g$ exists for $\alpha \geq 1/\pi$. For $\alpha > [\pi(1-\eta)^2]^{-1}$ this solution has $G = g < 1$, which means that agents do not all play pure strategies. When $\alpha \to [\pi(1-\eta)^2]^{-1}$, $G \to 1$ and the solution becomes independent of η. In other words, the solution merges with the solution for $\eta = 1$. In its turn this solution breaks down, with $\chi \to \infty$ and $H_1/N = \sigma^2/N \to 0$ when $\alpha \to 1/\pi$. Below this point, only solutions with $G < g$ and $H_1/N = 0$ exist. This behavior is well documented in figure 4. However, for $\eta > 0$, one needs to go beyond the simple approximation for $G_{a,b}$ and $r_{a,b}$ in Eq. (A.7). Therefore we shall refrain from a more detailed

[24] Note indeed that g has the interpretation of the overlap between two replicas of the same system, so that $g < G$ means that the two replicas are not identical.

discussion and rather refer the interested reader to a forthcoming publication [23].

References

[1] Arthur W. B., *Inductive reasoning and bounded rationality*, Am. Econ. Assoc. Papers and Proc **84**, 406, 1994.

[2] Erev I. and A.E. Roth, *Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria*, Am. Ec. Rev. **88**, 848 (1998).

[3] D. Fudenberg and D. K. Levine, *The theory of learning in games* (MIT Press, 1998).

[4] A. Rustichini, *Sophisticated players and sophisticated agents*, CentER discussion paper (October 1998).

[5] Casti J. L. *Seeing the light at El Farol*, Complexity (**1**)5, 7 (1996).

[6] Challet D. and Zhang Y.-C., *Emergence of cooperation and organization in an evolutionary game*, Physica A **246**, 407 (1997).

[7] Y.-C. Zhang, *Evolving models of financial markets*, Europhys. News **29**, 51 (1998).

[8] Savit R., Manuca R., and Riolo R., *Adaptative competition, market efficiency and phase transition*, Phys. Rev. Lett., **82**, 2203 (1990).

[9] Johnson N.F. et al.,*Volatility and Agent Adaptability in a Self-Organizing Market*, Physica A (**258**), 230 (1998).

[10] A. Cavagna, Irrelevance of Memory in the Minority Game, Phys. Rev. E **59**, R3783 (1999), preprint cond-mat/9812215

[11] Challet D. and M. Marsili, *Phase Transition and Symmetry Breaking in the Minority Game*, Phys. Rev. **E 60** (Dec. 1999).

[12] For a collection of papers and preprints on the minority game see the web site http://www.unifr.ch/econophysics/.

[13] Challet D. and Zhang Y.-C., *On the Minority Game : analytical and numerical studies*, Physica A **256**, 514 (1998).

[14] Johnson N.F. et al., *Crowds Effects and Volatility in a Competitive Market*, Physica A **269**, (1999).

[15] Challet D., M. Marsili and R. Zecchina, *Statistical Mechanics of Systems with Heterogeneous Agents*, preprint cond-mat/9904392 (1999).

[16] M. Mezard, G. Parisi, M. A. Virasoro, *Spin glass theory and beyond* (World Scientific, 1987).

[17] Challet D., M. Marsili and Y.-C. Zhang *Modeling Market Mechanisms*, Physica A **276**, 284 (2000).

[18] Myerson, R.B., *Game Theory: analysis of conflict*, Harvard Univ. Press (1997).

[19] Johnson N.F. et al., *Minority Game with Arbitrary Cutoffs*, preprint cond-mat/9903228 (1999).

[20] M. Woodford, *Learning to believe in sun-spots*, Econometrica **58**, 277 (1990).

[21] R. J. Aumann, *Rationality and bounded rationality*, in *Cooperation: Game-Theoretic approaches*, S. Hart and A. Mas-Colell Eds., NATO ASI Serie F, Vol. **155** (Springer, 1997).

[22] Weibull, J.W., (1995), *Evolutionary Game Theory*, The MIT Press, (Cambridge, Massachusetts; London, UK.

[23] A. De Martino, M. Marsili and R. Zecchina forthcoming.

[24] A. Rustichini, *Optimal properties of Stimulus-Response Learning Models*, Games and Economic Behavior (forthcoming, 1999).

[25] Luce R.D., *Individual Choice Behavior: A Theoretical Analysis*, New York, Wiley (1959).

[26] McKelvey, R.D. and T.R. Palfrey, *Quantal response equilibria for normal form games*, Games Econ. Behav. **10**, 6 (1995).

[27] A. Cavagna, J.P. Garrahan, I. Giardina, D. Sherrington, *A thermal model for adaptive competition in a market*, Phys. Rev. Lett. **83**, 4429 (1999).

[28] D. Challet and M. Marsili, *"Univeraslity" in the minority game*, forthcoming.

[29] Franke R., *Reinforcement learning in the El Farol problem*, preprint (March 1999).

[30] Hardin, G., *Science* **162**, 1243 (1968).

[31] S. W. Golom, *Shift Register Sequences*, (Aegean Park, 1981). We thank P. Flajolet for introducing us to this issue.

[32] Challet D. and M. Marsili, *Diffusion on random De Bruijn graphs*, forthcoming.

[33] G. Caldarelli, M. Marsili and Y.-C. Zhang, *A Prototype Model for Stock Exchange Markets*, Europhys. Lett. **40**, 479 (1997).